



The Bayes Limit

SUM BUNNY

The Bayes Limit

Learning, Drift, and the Missing Law of Stability

By Sum Bunny

2026



OPEN THROTTLE
CC BY 4.0

YOU ARE FREE

to share and adapt the material for any purpose,
even commercially, as long as appropriate credit is given.

THE BAYES LIMIT

LEARNING, DRIFT, AND THE MISSING LAW OF STABILITY

By Sum Bunny

I. The Minister and the Machine

In 1761,
Thomas Bayes
died leaving behind an unfinished manuscript on a narrow question:

How should belief change when new evidence arrives?

His friend
Richard Price
recognized its importance and presented it to the
Royal Society.

The idea did not transform the world immediately.

It waited.

For nearly two centuries, Bayesian reasoning remained mathematically elegant but computationally impractical. Only with the arrival of modern computing did it become operational at scale.

What followed was not a single invention, but a shift in method:

Systems could now update beliefs continuously under uncertainty.

II. The Inversion

Bayes introduced a reversal in reasoning:

From:

- given a system $\rightarrow$ predict outcomes

To:

- given outcomes $\rightarrow$ infer the system

$$P(H \mid D) = (P(D \mid H) \times P(H)) / P(D)$$

This rule does not produce certainty.

It produces **revision**.

A Bayesian system maintains a distribution over possible hypotheses and updates that distribution as new data arrives. In practice, many real systems compress this distribution into actionable estimates—working baselines that guide decisions.

III. From Inference to Systems

Modern technologies (from probabilistic classifiers to large language models and generative AI systems) operate within this general paradigm:

- maintain internal representations of uncertainty
- update those representations under new data
- act based on the most probable interpretations

This is not a claim that all modern AI is strictly Bayesian in implementation.

It is a claim that:

Probabilistic updating under uncertainty has become a dominant operational logic.

IV. What Bayes Does—and Does Not Do

Bayesian inference answers a specific question:

How should beliefs update given evidence?

It does not answer:

- whether updates converge
- whether systems remain stable under repeated updating
- how new hypotheses are generated
- how multiple updating systems coordinate
- how updating interacts with environmental limits

Bayesian inference operates over a defined hypothesis space. It does not, by itself, generate new categories or representations.

When systems expand the space of possible hypotheses, they are no longer performing inference alone—they are performing **generation**.

V. The Practical Assumption

In real-world applications, systems are often built as if the following were sufficient:

- continuous updating
- sufficient data
- improved estimation

Implicitly, this assumes:

Better updating leads to better system behavior.

This assumption can hold in stable, well-bounded environments.

It breaks in systems that are generative, adaptive, or distributed across many agents.

VI. Return Without Creation

Within the TUBE framework, Bayesian updating can be interpreted as a form of adaptive return.

From *TUBE*:

Systems stabilize by returning toward an internal baseline that can shift over time.

In this interpretation:

- the system maintains internal expectations
- incoming data creates deviation
- updates move the system toward revised expectations

This alignment is functional rather than exact. Bayesian systems typically operate over distributions, while TUBE expresses dynamics in terms of deviation and return.

In dynamical systems terms, the internal baseline functions as a **dynamic attractor**—a state toward which the system tends under correction. Adaptive return can therefore be understood as motion within a shifting basin of attraction, where stability depends on the system's ability to remain within that basin as it evolves.

VII. The Creation Problem

Real systems do more than update.

They:

- introduce new hypotheses
- expand representational space

- generate novel outputs

This is not updating.

It is **generation**—the expansion of the hypothesis space itself.

Generation changes the conditions under which inference operates.

It introduces a second requirement:

The system must remain coherent while producing novelty.

VIII. The Stability Constraint

This leads to a necessary condition for system viability:

The rate at which a system generates variation must not exceed its capacity to evaluate, correct, and integrate that variation.

In TUBE, this is expressed as:

$$D \leq R$$

Where:

- **D** = generative activity (hypothesis expansion, output novelty, structural change)
- **R** = return capacity (error correction, feedback fidelity, validation bandwidth, recovery mechanisms)

This is not a claim about correctness.

It is a constraint on **maintainability**.

IX. Operational Meaning

The constraint can be expressed in terms of rates:

- **D**: rate of generated variation
- **R**: rate and fidelity of evaluative feedback

In some domains, these admit information-theoretic approximations:

- **D** as an entropy rate of outputs
- **R** as the effective capacity of the feedback channel

When the rate of generated variation exceeds evaluative capacity:

Error is no longer reducible. It accumulates.

Not all feedback constitutes return. For R to function as a stabilizing force, it must preserve or recover a meaningful state relative to the system's prior structure. Feedback that is noisy, lossy, or uncorrelated with the system's baseline does not reduce error—it introduces additional variation. In such cases, apparent “correction” becomes indistinguishable from further generation.

R, therefore, is not the presence of feedback, but the capacity for **non-lossy, structure-preserving correction**.

X. Drift and Breakdown

When D exceeds R, systems do not become more creative.

They become unstable.

Examples across domains include:

- generative models producing ungrounded outputs
- feedback loops amplifying error faster than correction
- institutions expanding policy faster than evaluation capacity
- social systems generating narratives faster than they can be reconciled

In each case:

Generation outpaces correction.

This is not a failure of inference.

It is a failure of balance.

Drift is what happens when the system can no longer distinguish correction from generation.

XI. Beyond the Individual System

Bayesian inference operates primarily at the level of individual models.

Real systems exist across multiple layers:

- **R₂ (adaptive return):** individual updating
- **R₃ (coordination):** interaction among agents
- **R₄ (constraint):** environmental and material limits

Even when individual systems update correctly, the larger system can fail:

- through misalignment between agents
- through insufficient coordination

- through violation of environmental constraints
- Stability is not a property of inference alone.

It is a property of **coordinated inference under constraint**.

XII. The Bayes Limit

Bayesian reasoning remains a powerful method for inference under uncertainty.

But it operates within a boundary:

It governs how beliefs change, not whether systems remain viable under those changes.

This boundary can be stated simply:

A system can update correctly and still fail.

This is the **Bayes Limit**.

XIII. Continuity and Identity

For a system to remain stable, it must preserve continuity with its prior state.

Return capacity is not only about correcting error.

It is what allows a system to remain identifiable across time.

Without sufficient return:

- updates do not refine the system
- they replace it

Drift is not just error.

It is **loss of identity under change**.

XIV. Final Statement

Bayes provided a rule for learning from evidence.

Modern systems have scaled that rule across machines, institutions, and networks.

What has not scaled at the same rate is the capacity to stabilize what is learned.

Bayesian inference governs how systems update beliefs within a given space.

The Bayes Limit defines the boundary beyond which those updates can no longer be stably maintained.

Learning without constraint does not produce intelligence.

It produces drift.

Appendix: Transmission

(Non-Analytical Layer)

First edition published on the Internet Archive, 2026.
Part of the Sum Bunny Canon (Ark Hive Press).

This work was developed with AI-assisted analysis tools. All arguments, structures, and final judgments are the author's.

Hey Josiane,
Happy Mayday Eve!
Sum Bunny Loves You!!!

T. Bayes.

FREE FLASH